between the headword and a concordance term. As such, the presence of concordance terms for a headword does not augment the semantic information for that headword, nor does it facilitate any NL processing task, such as resolving syntactic ambiguity. Specifically, Amsler only explored discerning taxonomic information (*hypernym, hyponym*) from the concordances, which are of only limited use in NL processing. Amsler's concept was also cited by Chodorow et al. (1985) in developing a tool for helping human users disambiguate hyper/hyponym links among pairs of lexical items; again, however, this approach was limited to *hypernym* and *hypernym_of* semantic relations. Boguraev et al. (1989) discuss *distributed lexical knowledge*, in which the structure of each lexical entry is represented explicitly and the dictionary as a whole can be queried using a strategy of either query-by-example or unification. However, in Boguraev et al. (1989), the information that can be queried is only that which can be conveyed by the structure of a lexical entry, not the contents of either its definition and/or example sentences (they describe no method for extracting semantic information from the contents of the definition and/or example sentences). Moreover, as described in Boguraev et al. (1989), the distributed lexical knowledge can only be discovered by constructing queries manually, and it is described to be useful for the researcher who wants to acquire lexical information; this is in contrast to our system, which typically constructs a semantic knowledge base for consumption by a computer application, and is only incidentally useful to a researcher, and this knowledge base contains all of the relevant information pertaining to a lexical entry on that entry itself, so no query mechanism is required to associate the information which is found distributed in the on-line dictionary.

In accordance with the preferred embodiment of the present invention, a lexical knowledge base is compiled automatically from a machine-readable source, such as an on-line dictionary or unstructured text, obviating many of the drawbacks associated with the foregoing prior art techniques. The preferred embodiment of the invention makes use of "backward linking" by which inverse semantic relations are discerned and used to augment the knowledge obtained

from traditional forward-linking analysis of the parsed text. Iteration of this technique can further enhance the results.

The foregoing and additional features and advantages of the present invention will be more readily apparent from the following detailed description thereof, which proceeds with reference to the accompanying drawings.

Brief Description of the Drawings

Fig. 1 is a flow chart illustrating a method for generating a lexical knowledge base according to one embodiment of the present invention.

Fig. 2 is a logical form produced from a definition of the word "market" in accordance with one embodiment of the present invention.

Fig. 3 is a semantic relation structure corresponding to the logical form of Fig. 2.

Fig. 4 is an inverted semantic relation structure derived from the semantic relation structure of Fig. 3.

Fig. 5 is an illustration of an exemplary semantic network discussed in the specification.

Fig. 6 shows the semantic network of Fig. 6 with elaboration.

Fig. 7 is an illustration of another exemplary semantic network.

Detailed Description

We now describe our use of DB methods to automatically create a semantic knowledge base from an on-line dictionary. Our approach builds on the work of Jensen and Binot (1987) and